

FINAL REPORT

Signal Processing and Modeling for UXO Detection and Discrimination in Highly Contaminated Sites

SERDP Project MM-1281

FEBRUARY 2007

Leslie Collins
Duke University

This document has been approved for public release.



Strategic Environmental Research and
Development Program

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE FEB 2007		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Signal Processing and Modeling for UXO Detection and Discrimination in Highly Contaminated Sites				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Duke University				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 61	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

This report was prepared under contract to the Department of Defense Strategic Environmental Research and Development Program (SERDP). The publication of this report does not indicate endorsement by the Department of Defense, nor should the contents be construed as reflecting the official policy or position of the Department of Defense. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the Department of Defense.

Table of Contents

List of Acronyms:	ii
List of Figures:	iii
Acknowledgements:	viii
Executive Summary:	ix
Background:	1
Objective:	1
Materials and Methods:	1
Results and Accomplishments:	2
Conclusions:	48
Technology Transfer:	49
Manuscripts Published and Presentations Given:	49

Acronyms

AIC – Akaike information criterion
BOR – body of revolution
BSS – blind source separation
EDA – eigenvalue decomposition analysis
EMI – electromagnetic induction
EM-61 – Geonics’ four gate time-domain EMI system
EM-63 – Geonics’ multi-gate time domain EMI system
FEM – Finite Element Model
GEM-3 – Geophex’s frequency-domain EMI system
GLRT – Generalized Likelihood Ratio Test
H0 – Hypothesis 0, or null hypothesis, or non-UXO hypothesis
H1 – Hypothesis 1, or target/UXO hypothesis
ICA – independent components analysis
IEEE – Institute of Electrical and Electronic Engineers
IPR – In-progress review
MDL – minimum description length
NRL – Naval Research Laboratory
ROC – receiver operating characteristic
SERDP – Strategic Environmental Research and Development Program
SNR – signal to noise ratio
TRGS – Transactions on Geoscience and Remote Sensing
UXO – Unexploded ordnance.
WES – Waterways Experiment Station

List of Figures

Figure 1. Schematic of a conducting cylinder and ring, the two used to constitute a composite target. Also denoted is the coordinate system used to define the angle of observation θ . In the computations considered here the cylinder length and diameter are 2.54 cm; the ring is defined by an inner diameter of 2.6 cm, outer diameter 3.9 cm and thickness 5 mm; the distance $d=5$ mm; and the cylinder and ring conductivities are $\sigma=3 \times 10^7$ S/m and $\sigma=5 \times 10^6$ S/m, respectively.

Figure 2. Model predictions for two different object orientations, with a ring shown on the left and a cylinder shown on the right. Predicted voltage is plotted as a function of frequency.

Figure 3. EMI response of the composite ring-cylinder target in Figure 1, as observed at angle of observation 0° (top left) 180° (top right) and 90° (bottom). The squares denote the results of the direct FEM solution and the circles the model fit based on a dipole representation of each target part. The data used to produce these curves were computed at the frequencies associated with the points, denoted by squares and circles.

Figure 4 . Model fits and measured data for two different object orientations of an 81mm mortar. A) zero degrees, B) 180 degrees

Figure. 5. Model fits and measured data for ninety degree orientation of an 81mm mortar.

Figure. 6. Model predictions and measured data for an 81mm mortar at 45 degrees. Prediction is based on model fits from zero, one eighty, and ninety degrees

Figure 7. Single dipole fit to measured GEM data from a 105 mm UXO (left) and two-dipole fit (right) to the same data.

Figure 8. Confusion matrices for 4-target classification problem. True targets labeled in rows, classified targets labeled in columns. Classification results for single dipole model on the left, two dipole model on the right.

Figure 9. Average percent classification obtained from the confusion matrices obtained for both one- and two-dipole fits. Three cases are considered depending on whether additive white Gaussian noise is present, and how much variability/uncertainty there is in the target moment parameters. Generally, performance is stable under reasonable amounts of noise, but uncertainty in the target parameters degrades performance. Higher levels of noise also degrade performance. Performance changes in the single dipole model are not statistically significant.

Figure 10. ROCs showing performance of an error-based discriminator for determining between the multiple dipole/single dipole hypotheses where ROCs are parameterized by the distance separating the two objects. Simulation considers EM 61 data.

Figure 11. ROCs showing performance of an error-based discriminator (left) and GLRT-based discriminator where H_0 data is obtained from training (right) for determining between the multiple dipole/single dipole hypotheses. Simulation considers EM 61 data.

Figure 12. AIC output (left) and MDL output (right) for a single object. Minimum shown with red square.

Figure 13. AIC output (left) and MDL output (right) for two objects. Minimum shown with red square.

Figure 14. GEM response at 2790 Hz (red = in-phase, blue=quadrature) for two objects individually (circle and plus sign), when summed (dashed) and measured with both objects present (solid). Objects are 9 inches apart and 4 inches from the sensor (top) or 9 inches apart and 7 inches from the sensor (bottom).

Figure 15. Left panel: Probability of correct classification of 5 objects that appear in isolation as a function of SNR (simulated data). Right panel: Post-ICA classification of 5 objects that appear in pairs as a function of SNR (simulated data). Object identifications are: 40 mm (object 1), 60 mm (object 2), 81 mm (object 3), M42 submunition (object 4), and a 155 mm (object 5).

Figure 16. Confusion matrices for classifier where objects occur in isolation (top) or in pairs (bottom). Average percent correct for objects in isolation is 95%, and is 85% when occur in pairs when ICA is used, 14% if ICA is not used.

Figure 17. Original and extracted sources obtained for two (M42 and 40mm) of the objects considered in the source separation experiment. Left panels are for the objects oriented vertically, right panel for the objects oriented transverse. The two objects were measured together and separated by 1". Red curves show original sources, blue are the sources extracted using EDA. Solid lines plot in-phase data, dashed lines plot quadrature data.

Figure 18. Confusion matrices listing fractional percent correct identification for experimental data when two targets are present simultaneously in the field of view of the sensor. True targets labeled in rows, classified targets labeled in columns.

Figure 19. Example of experimental setup for UXO + multiple clutter experiment.

Figure 20. Portion of the experimental results from the UXO + multiple clutter experiment. Targets present are shown in the first column. Classification results for the UXO signature extracted shown in the 2nd and 3rd columns. A 60 mm was always present (target 1).

Figure 21. Example of the simulation setup for the circumscription problem. Brown area contains no UXO, green area contains widely separated anomalies and dark grey area contains overlapping objects. Cyan lines show simulated transects, data from which are provided to the various algorithms to determine the number of objects present as a function of distance along the transect.

Figure 22. Example results from a single simulation, a single transect. Signal complexity is plotted as a function of distance along the transect. Red curve is a smoothed version of the output.

Figure 23. A comparison of the performance of classification techniques with and without incorporation of ICA. Classifiers are trained on calibration measurements of isolated UXO. Performance is measured in terms of area under the ROC for which 1 indicates perfect performance and 0.5 indicates random chance.

Figure 24. ROC for each of the four measured objects. Note that the 60 mm object achieves perfect detection at zero probability of false alarm. Only one source is extracted by ICA per object pair.

Figure 25. ROC for each of the four measured objects. Two sources are extracted by ICA per object pair. 60 mm curve follows 81 mm curve where color is ambiguous.

Figure 26. A comparison of the performance of ICA with the original, processed data and the original data with additional pre-processing. Performance is measured in terms of area under the ROC for which 1 indicates perfect performance and 0.5 indicates random chance.

Figure 27. The four signals used for the toy problem simulation. All correlation coefficients between signals were less than 0.05.

Figure 28. An example of the different levels of overlap for the toy problem simulation. The color bars indicate amplitude, where blue equals zero and red equals 1. The top portion of each subplot represents an object that remains in the same position, while the bottom portion of each subplot represents an object that is brought closer to the first object such that their spatial responses overlap (i.e. their signatures are mixed).

Figure 29. Result of increasing the spacing between samples for the UXO model simulation for 1-5 clutter objects. The performance, measured in area under the ROC, was averaged across separation distance.

Figure 30. Results for the case in which clutter no longer has a constant separation distance from the target, i.e. clutter is randomly placed within the square meter without relation to target location. Spacing is increased between samples for the UXO model simulation for 1-5 clutter objects.

Figure 31. Result of increasing the spacing between samples for the UXO model simulation for 1-5 clutter objects. In this case, an energy detector rather than a generalized Gaussian detector (Fig. 29) was used. The performance, measured as area under the ROC, was averaged across separation distance.

Figure 32. Comparison of performance for the UXO model simulation for three different methods of sampling: asterisk pattern, energy partitioning, and energy threshold. Performance was measured in terms of area under the ROC.

Figure 33. The effect of varying the number of samples used for the toy problem simulation. In this case, 0% overlap was used, and the amplitude patterns were modeled from the dipole model, with random orientations. The N highest energy samples were used where N ranges from 800 to 25. Performance was measured in terms of correlation coefficients.

Figure 34. The effect of varying the number of samples used for the toy problem simulation. In this case, 0% overlap was used, and the amplitude patterns were modeled from the dipole model, with random orientations. The N highest energy samples were used where N ranges from 800 to 25. Signal 4 (the square wave) was replaced with another window pulse train that had a correlation coefficient with signal 1 of 0.3. Performance was measured in terms of correlation coefficients.

Figure 35. Example of two amplitude patterns used with the toy problem simulation: simplistic peak with one sample different, and random amplitude patterns.

Figure 36. Average correlation coefficients for each signal for the toy problem simulation. The two types of amplitude patterns depicted in Fig. 36 are compared.

Figure 37. Comparison of the performance for different types of amplitude patterns for the toy problem simulation. Dotted lines represent both the random amplitude patterns (filled symbols) and one-sample-different peak amplitude patterns (hollow symbols). The solid lines represent the model-based amplitude patterns.

Figure 38. Comparison of the variability for different types of amplitude patterns for the toy problem simulation. Dotted lines represent both the random amplitude patterns (filled symbols) and one-sample-different peak amplitude patterns (hollow symbols). The solid lines represent the model-based amplitude patterns.

Figure 39. The average of the amplitude patterns producing the 200 highest correlation coefficients for each signal.

Figure 40. The average of the amplitude patterns producing the 200 lowest correlation coefficients for each signal.

Figure 41. Simplistic amplitude patterns generated for the toy problem simulation: peak, ring, 2 peaks, 2 peaks flipped.

Figure 42. Comparison of performance and variance for two cases of the toy problem simulation: randomly choosing amplitude patterns from a set of four and using random amplitude patterns. Error bars represent one standard deviation from the mean.

Figure 43. The effect of the number of sources extracted by ICA for the toy problem simulation. In this case, 0% overlap was used, and the amplitude patterns were modeled from the dipole model, with random orientations. The number of sources varied from 1-7, with 4 (the number of signals) being the correct value. Performance was measured in terms of correlation coefficients.

Figure 44. UXO-UXO pair classification accuracy using ICA template method vs. different noise levels under the condition that two UXO present.

Figure 45. UXO-Clutter pair classification accuracy using ICA template method vs. different noise levels under the condition that UXO-Clutter pairs present – 2 UXO in library and in data.

Figure 46. UXO Classification Accuracy using ICA template method vs. different noise levels under the condition that UXO-Clutter pairs present – 4 UXO present in library and in data.

Figure 47. Comparison between iterative fit (left) and maximum likelihood fit (right). In each set of panels, the measured (field) data is shown on the top left, the two individual fits to the data are shown on the bottom left, and the synthesized data from the two individual fits is shown on the top right.

Figure 48. Area under the curve for classification of each ordnance type from the SERDP/NRL/WES database. AUC is plotted as a function of the number of grid points used to perform ICA. Training data for the classifier was obtained from the isolated UXO items.

Acknowledgements

The following individuals and organizations contributed to the success of this project: Dr. Lawrence Carin, Dr. Chandra Throckmorton, Dr. Stacy Tantum, Dr. Yan Zhang, Mr. Wei Hu, Ms. Chunmei Wang, Ms. Debbie Schofield, AETC, WES, NRL, and Geophex.

Executive Summary

In this effort, we considered the problem of classifying closely spaced UXO. When two UXO are in close proximity, their signatures as measured by electromagnetic induction (EMI) sensors co-mingle, and traditional classification algorithms cannot be utilized. We explored the use of independent components analysis (ICA), a technique from the blind source separation (BSS) literature as a pre-processing step by which to separate the individual UXO signatures from the mixtures measured by the EMI sensor. Simulations indicated that this procedure could succeed and restore some level of classification performance in the case of overlapping signatures. Test-stand data also indicated that the ICA-based classification approach showed promise. Finally, testing the algorithms on data collected by NRL/WES suggested that in some cases ICA could be used as a pre-processing step and that closely-spaced UXO could be classified.

During the course of this effort, several studies were performed to try to assess the limitations of the ICA-based approaches for discrimination of closely spaced objects. Correlation between object signatures can degrade performance, and alternate techniques based on ICA were considered as a remedy to this issue. Other issues including sampling density, noise susceptibility, and others were also investigated.

I. Background

In this project, we investigated the phenomenological aspects of the unexploded ordnance (UXO) detection, location, and discrimination problem with EMI sensors, with a focus on highly contaminated sites. In the majority of clearance scenarios, UXO co-mingle with extensive surface and subsurface clutter, such as metal pieces from shattered ordnance (e.g. fins). Moreover, near the “bull’s-eye” of a former range, many UXO are in close proximity, with the classification problem exacerbated by intermingling anthropic clutter. Furthermore, naturally occurring magnetic geologic noise often adds to the complexity of the discrimination task. Until recently detection algorithms could not distinguish between buried UXO and clutter, leading to many false alarms. Over the last several years, modern geophysical techniques have been developed, these merging more-sophisticated sensors, underlying physical models and statistical signal processing algorithms. Such approaches have yielded reduced false alarms. In particular, for sites where anomalies are well separated, it has been shown that the combination of phenomenological models and advanced signal processing can markedly decrease the time required to remediate a site by classifying UXO and non-UXO items correctly. For highly contaminated regions, however, the signatures of multiple anomalies often overlap, vitiating the utility of many of the newer techniques. We pursued a program employing a synergistic use of advanced phenomenological modeling and signal processing algorithms to address this problem. In the end, progress was made toward solving this problem, and issues that limited the effectiveness of our proposed solutions were highlighted.

II. Objective

The research program had two principal objectives: (1) the development of new physics-based signal-processing approaches applicable to scenarios in which responses from multiple UXO and clutter items co-exist in a sensor signal, with the goal of discrimination; and (2) the use of information-theoretic measures to define the types of scenarios for which UXO and clutter density is too high to reliably perform classification, necessitating a direct mechanical excavation of an entire region. The first objective was addressed by the parallel development of phenomenological models and statistical signal processing algorithms. The latter topic addressed circumscription of those regions, presumably in the vicinity of a former bull’s-eye, for which discrimination of individual UXO and clutter is intractable due to the high density of target/clutter overlap and the limited information in available sensor data. Implicitly, this latter objective could lead to tools that allow circumscription of those regions for which identification of individual UXO adds value.

III. Materials and Methods

We pursued a research program that focuses on phenomenological modeling of EMI responses to multiple objects, and on the development of the associated physics-based signal processing algorithms. We used the models to motivate the statistical algorithms,

and to generate a database of simulated signals on which we can test our algorithms. As we developed models capable of predicting the EMI signature of multiple UXO and anthropic-clutter items, we used these models to develop detection and discrimination performance bounds for the scenario of multiple proximate objects. We also investigated both traditional and novel signal processing approaches to allow both the detection of the existence of multiple objects, as well as the separation of the individual signatures from the cumulative signature. Following the model-based simulations, we transitioned to processing data measured from a field pit.

There are two principal sensors used currently in modern UXO sensing: electromagnetic induction (EMI) and magnetometers. Magnetometer and EMI models have been developed by Duke and others for conducting and ferrous targets in isolation, and it was therefore of interest to address how these models must be augmented to handle the multi-target problem. Initially, our models had been applied to simple, isolated UXO. There are many cases for which one would be interested in more-complicated targets. In the context of individual UXO, many ordnance are composed of multiple parts (body, rings, fins, etc.) each of which may contribute its own dipole response to the composite signature. We extended the above EMI resonant-dipole model to the case of targets with multiple parts, utilizing iterative techniques.

In terms of signal processing, we developed a three stage processing procedure. The first stage was used to determine whether multiple objects are present in the sensor's field of view, and the second was processed to determine the likelihood that the sensor data are associated with a UXO item. In agreement with reports from other labs [e.g. B. Barrow and H. H. Nelson, "Model-Based Characterization of Electromagnetic Induction Signatures Obtained with the MTADS Electromagnetic Array," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, pp. 1279-1285, 2001.; T. H. Bell, B. Barrow, and J. T. Miller, "Subsurface Discrimination Using Electromagnetic Induction Sensors," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, pp. 1286-1293, 2001.], our preliminary analyses indicated that utilization of a dipole-fit parameter to EMI data can be used as a pre-screener to determine locations where it is likely that multiple objects are present. Initial stages of the object separation algorithms also provided estimates of the number of objects that could be used as a prescreener. In addition to further investigating the use of Bayesian techniques we have used to successfully discriminate UXO objects from clutter in relatively uncontaminated sites, we also investigated independent component analysis (ICA) techniques for blind source separation (BSS). The ICA approach had not been applied to date for the subsurface object detection and identification problem. Our focus was on the development of signal processing algorithms that rigorously incorporate the underlying physics characteristic of the sensor and the anticipated UXO target. These algorithms could provide both a mechanism by which to detect UXO in highly contaminated environments, and also a performance bound which will define the point at which alternative clearance technologies should be employed. For the acquisition of field data on which we could more fully test the performance and robustness of our algorithms, we will rely on the sponsor as well as our connections with government laboratories (NRL) and industry (Geophex).

IV. Results and Accomplishments

Modeling

We extended the EMI magnetic-dipole model to the case of complex targets, such as UXO. We demonstrated that the simple magnetic-dipole model used in our previous modeling efforts can be extended readily to the case of complex targets, by considering *multiple* offset frequency-dependent magnetic dipoles, associated with various parts of the target. For example, in the context of a UXO, magnetic dipoles may be used to represent localized firing rings or fins on the ordnance, while distinct and spatially separated dipoles are used to represent the UXO's main body (or other components). The simple model is adequate for relatively simple targets with dimensions small relative to the target-sensor distance. For example, the model fits the EMI response for an object that is identical as observed from both ends of the target. This is true for a cylinder, for example, but not in general for a UXO. The formulation of the single dipole model that we have used is

$$\mathbf{M}(\omega) = \mathbf{z}\mathbf{z}[m_z(0) + \sum_k \frac{\omega m_{zk}}{\omega - j\omega_{zk}}] + (\mathbf{x}\mathbf{x} + \mathbf{y}\mathbf{y})[m_p(0) + \sum_i \frac{\omega m_{pi}}{\omega - j\omega_{pi}}]$$

where \mathbf{z} is a unit vector in the z direction, and \mathbf{x} and \mathbf{y} correspond to orthogonal unit vectors, each perpendicular to \mathbf{z} . The terms $m_z(0)$ and $m_p(0)$ account for the induced magnetization produced for ferrous targets (valid down to static magnetic-field excitation, $\omega \rightarrow 0$), and the terms in the summations account for the frequency-dependent character. For simple targets, typically we only require the first term in each sum, representative of the principal dipole mode along each of the principal axes. We extended this model to the form

$$\begin{aligned} \mathbf{M}(\omega) &= \sum_{n=1,N} \left\{ \mathbf{z}\mathbf{z}[m_z^n(0) + \sum_k \frac{\omega m_{zk}^n}{\omega - j\omega_{zk}^n}] + (\mathbf{x}\mathbf{x} + \mathbf{y}\mathbf{y})[m_p^n(0) + \sum_i \frac{\omega m_{pi}^n}{\omega - j\omega_{pi}^n}] \right\} \delta(\mathbf{r} - \mathbf{r}_n) \\ &= \sum_{n=1,N} \mathbf{M}_n(\omega) \delta(\mathbf{r} - \mathbf{r}_n) \end{aligned}$$

where \mathbf{r}_n represents the location of the n th set of dipoles. Note that we have assumed that the local coordinate system of each dipole set (\mathbf{x} , \mathbf{y} , \mathbf{z}) is the same, although this need not be true in general. Although the model appears to represent a significant escalation in complexity from the simple model, we note that in practice (for actual UXO) we typically only require a small number of terms N .

We considered a conducting cylinder of length 2.54 cm, diameter 2.54 cm and conductivity $\sigma = 3 \times 10^7$ S/m. Other simulations with different geometries were also considered, and similar conclusions were drawn. In addition, we considered a ring of inner radius 2.6 cm, outer radius 3.9 cm, thickness 5 mm, and conductivity $\sigma = 5 \times 10^6$ S/m. The targets are sensed via a simulated GEM-3 frequency-domain EMI sensor. The

cylinder and ring are used to realize a composite target, as depicted in Figure 1. The center of the sensor coils are positioned 12 cm from the center of the cylinder. Our objective is to consider the fit of the simple parametric model to the FEM-computed EMI signature of the cylinder and ring alone. The FEM model was developed under previous SERDP support. We then consider the accuracy of the composite model in the context of the cylinder-ring composite target. In this test we first perform the model fit of the isolated targets, from which we extract the parameters required in the single dipole model. These same parameters are then used in the composite model, but now the associated dipoles are positioned at the center of the cylinder and ring respectively (i.e. the magnetic dipoles associated with the ring and cylinder are offset with respect to each other, to reflect the respective positions of the corresponding target parts).




Figure 1. Schematic of a conducting cylinder and ring, the two used to constitute a composite target. Also denoted is the coordinate system used to define the angle of observation θ . In the computations considered here the cylinder length and diameter are 2.54 cm; the ring is defined by an inner diameter of 2.6 cm, outer diameter 3.9 cm and thickness 5 mm; the distance $d=5$ mm; and the cylinder and ring conductivities are $\sigma=3 \times 10^7$ S/m and $\sigma=5 \times 10^6$ S/m, respectively.

In Figure 2 we depict the in-phase and quadrature components of the computed voltage computed for the sensor, with the ring (left) and cylinder (right) in isolation. Results are shown with the coil axes aligned with the target axes (0°) and orthogonal to the target axes (90°). The solid curves represent the FEM computations and the points the single-dipole model fits. In Figure 3 we compare the rigorous, *coupled* FEM solution of the ring-cylinder composite, *vis-à-vis* the multi-dipole model fit, using the model parameters extracted from the *individual* targets. In Figure 3 we show comparisons between the FEM and parametric models for angles 0° , 90° and 180° . The parametric model, which here ignores coupling between the components, is generally in good agreement with the rigorous FEM computations, particularly for 90° . It is also important to note that the composite target appears different to the sensor as observed at all angles, due to the fact

that the ring is not in the center of the cylinder. A single set of dipoles, as in the simple model, does not capture these differences.




Figure 2. Model predictions for two different object orientations, with a ring shown on the left and a cylinder shown on the right. Predicted voltage is plotted as a function of frequency.

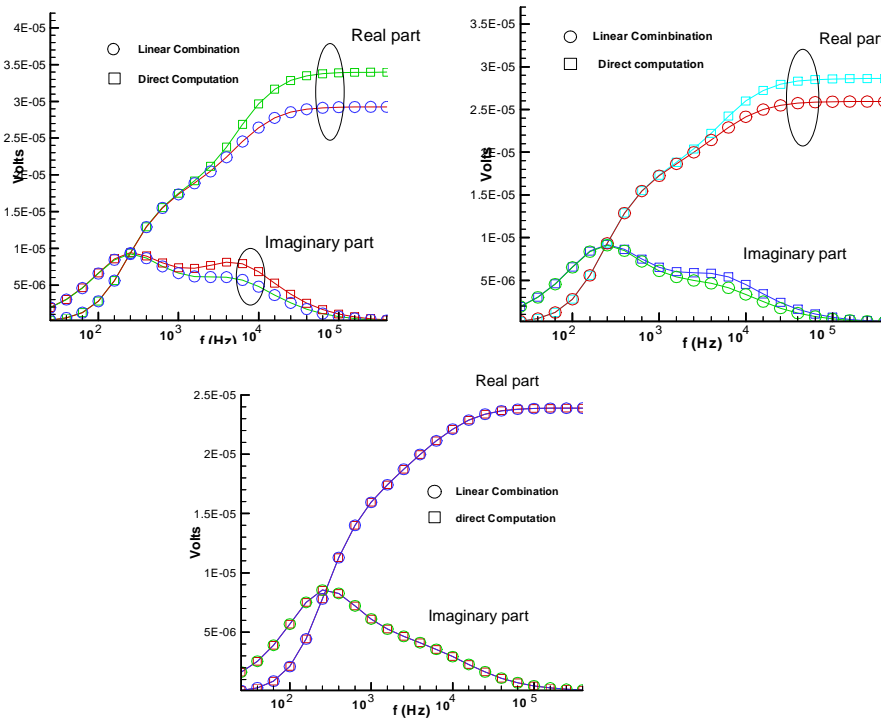


Figure 3. EMI response of the composite ring-cylinder target in Figure 1, as observed at angle of observation 0° (top left) 180° (top right) and 90° (bottom). The squares denote the results of the direct FEM solution and the circles the model fit based on a dipole representation of each target part. The data used to produce these curves were computed at the frequencies associated with the points, denoted by squares and circles.




Figure 4 . Model fits and measured data for two different object orientations of an 81mm mortar. A) zero degrees, B) 180 degrees




Figure. 5. Model fits and measured data for ninety degree orientation of an 81mm mortar.

We also considered this model in the context of *measured* GEM-3 data, for an actual ordnance, an 81 mm mortar. In Figures 4 and 5 we present measured data (points) and the EMI fit based on the multi-dipole, for sensor angles of observation 0° and 180° . Note the significant differences in the EMI signature as viewed along its axis from the top (0°) and from the bottom (180°), this motivating the composite model. As a more-challenging

test of the model, we utilize the parameters extracted in the context of Figures 4 and 5 – based on frequency-domain observations at three angles – to *predict* the complex frequency response at a fourth angle not observed when performing the model fit. Results are shown in Figure 6 for GEM-3 measured data at 45° , with a comparison to the multi-dipole model based on parameters extracted from data at three separate observation angles. We see in Figure 6 an encouraging comparison between the parametric model and the measured data. The results in Figures 4-6 also demonstrate the generally strong aspect dependence to the frequency-domain EMI signatures of actual ordnance. However, it is important to emphasize that although the results in Figures 4-6 show significant variation with orientation, each example is characterized by the same magnetization tensor \mathbf{M} , with the aspect-dependence of the signature characterized by \mathbf{U} . This implies that \mathbf{M} captures the underlying structure of the target itself, independent of orientation, and consequently the parameters of \mathbf{M} (not \mathbf{U}) are used in the classifiers.




Figure 6. Model predictions and measured data for an 81mm mortar at 45 degrees. Prediction is based on model fits from zero, one eighty, and ninety degrees




Figure 7. Single dipole fit to measured GEM data from a 105 mm UXO (left) and two-dipole fit (right) to the same data.

Above, we demonstrated that the two-dipole model provided better fits to the FEM model data. We also showed that a two-dipole model predicted experimental data that

had not been used in the fitting algorithm better than a single dipole model for an 81 mm mortar. Next, we considered additional UXO items and showed that the multiple-dipole fit was better than a single dipole fit for all complex ordnance considered, and did not degrade the fit for simpler ordnance items (results for 37mm, 60mm, and 105 mm ordnance items reported during the May 2003 IPR). Figure 7 shows an example comparison for a single dipole fit (left) and 2 dipole fit (right) for a 105 mm at 0 degrees. Similar results were seen for all complex ordnance at all orientations. A complete data set is available upon request, but these results indicated that the multi-dipole model does indeed improve fit error for complex ordnance items. The tradeoff, of course, is the number of parameters that must be estimated from the data.

In addition, we considered whether or not these improved fits actually resulted in improved discrimination performance. We considered four targets, the first three of which were “complex” and consisted of two dipoles at different locations. The stronger of the dipoles in all four cases was the same. The fourth target was “simple”, consisting of only a single dipole. Both one and two dipole fits were computed and used in a Bayesian classifier. To render the problem more difficult, the stronger dipole for all targets (the only dipole for the fourth target) was identical. Classification performance was computed as a function of noise level and for both deterministic target parameters and random target parameters. The random parameter case was considered in order to assess the impact of ordnance variability as suggested following the May 2003 IPR. An example confusion matrix obtained for classification with the single dipole model (left) and two dipole model (right) is shown in Figure 8. The average percent correct classification improves by almost 50%, indicating that the better fits associated with the two dipole model do positively impact discrimination performance. This trend is observed for all noise levels, and under conditions of parameter uncertainty. A few of the results for this analysis are summarized in Figure 9. The complete data set is available upon request.

Results with single dipole

Classified Target

T1

T2

T3

T4

True Target

T1

T2

T3

T4

0

0.16

0

0.84

0

0.66

0

0.34

0

0

0.4

0.6

0

0

0

1

AV % Correct = 51.4

Results with two dipoles

Classified Target

T1

T2

T3

T4

True Target

T1

T2

T3

T4

1

0

0

0

0

0.98

0

0.02

0

0

1

0

0

0

0

1

AV % Correct = 99.5

Figure 8. Confusion matrices for 4-target classification problem. True targets labeled in rows, classified targets labeled in columns. Classification results for single dipole model on the left, two dipole model on the right.

Case	Average Percent Correct	
	One Dipole	Two Dipole
No Noise, No Uncertainty	51.5	99.5
20 db AWGN, No Uncertainty	52	99.5
20 db AWGN, 10% Uncertainty	53.5	92

Figure 9. Average percent classification obtained from the confusion matrices obtained for both one- and two-dipole fits. Three cases are considered depending on whether additive white Gaussian noise is present, and how much variability/uncertainty there is in the target moment parameters. Generally, performance is stable under reasonable amounts of noise, but uncertainty in the target parameters degrades performance. Higher levels of noise also degrade performance. Performance changes in the single dipole model are not statistically significant.

One issue that has yet to be addressed is how to discriminate between a single UXO that is complex enough to need to be modeled with multiple dipoles and multiple single-dipole objects. While this question has not been addressed in this research, we have considered possible approaches to addressing this question. In general, the question we are trying to address is whether there are multiple dipoles in the field of view or not. At that point, we could then pose the question of whether the data could come from a UXO with a known multi-dipole model, a UXO with a single dipole model, or whether the dipoles are not associated with UXO. This complexity will be considered in future efforts.

Finally, the modeling assumes a BOR since we are most interested in considering UXO from the theoretical standpoint. For field work, we will not be utilizing the BOR assumption and thus will be able to ascertain whether an object's features extracted from a dipole model are indicative of a BOR or not.

Signal Processing

Prescreener

The first task was development of a pre-screener that can determine when multiple objects are present. Initially, we considered a pre-screener that detects the presence of multiple dipoles in spatially collected data. To determine whether multiple dipoles were present in simulated data, we utilized a simple Bayesian processor to decide between the hypothesis that two objects were present versus the hypothesis that one object was present. The simulated objects were single dipoles under H_0 and multiple dipoles under H_1 . Their moments were based on moment estimates from 60 mm and 81mm test data. The decision statistic that was utilized was the goodness of fit metric (error) between the simulated EMI data and the signature predicted by a single dipole model. For multi-frequency EMI data, such a simple prescreener provided excellent performance for almost all object separation distances. Figure 10 shows the ROC performance for a

variety of object separations for an EM 61 data, in which objects had to be considerably more distant to achieve performance similar to that for the multi-frequency system considered previously. In this curve, the probability of detection implies probability of detecting two objects when two are present versus the probability of false alarm. The latter is the probability of deciding that two objects are present when only one is present.

Because this prescreener was not as effective for EM 61 data, alternatives were considered. We considered a GLRT processor to test the hypotheses H_1 : more than one dipole present and H_0 : one dipole present. The data used by the processor was the extracted moments and moment ratios obtained using a single dipole fit. We evaluated performance of this GLRT under a variety of assumptions regarding the training data: specifically whether the density functions for the data under both H_0 and H_1 were available, or whether only H_0 prior information was available (as would most likely be true on a test site). Our results indicated that for many different combinations of UXO-like objects and clutter, satisfactory prescreener performance could be obtained when only H_0 training data was known. Figure 11 shows an example in which the fitting error metric that was used for the GEM-type prescreener was very ineffective (left), whereas the GLRT processor evaluated using training data for only H_0 was substantially better (right).




Figure 10. ROCs showing performance of an error-based discriminator for determining between the multiple dipole/single dipole hypotheses where ROCs are parameterized by the distance separating the two objects. Simulation considers EM 61 data.




Figure 11. ROCs showing performance of an error-based discriminator (left) and GLRT-based discriminator where H_0 data is obtained from training (right) for determining between the multiple dipole/single dipole hypotheses. Simulation considers EM 61 data.

We also considered two information-theoretic criteria that could be applied to frequency-domain data to estimate the number of dipoles: the Akaike Information Criterion (AIC) and a Minimum Description Length (MDL) criterion. Given a set of eigenvalues of the estimated covariance matrix, and k frequency samples, the AIC and MDL are given by

$$AIC(i) = -2(m-i)k \ln \frac{g(\lambda_{i+1}, \dots, \lambda_m)}{a(\lambda_{i+1}, \dots, \lambda_m)} + 2i(2m-i)$$

$$MDL(i) = -(m-i)k \ln \frac{g(\lambda_{i+1}, \dots, \lambda_m)}{a(\lambda_{i+1}, \dots, \lambda_m)} + \frac{1}{2}i(2m-i) \ln k$$

where $g(\bullet)$ and $a(\bullet)$ denote the geometric and arithmetic mean of their arguments,

$$g(\lambda_{i+1}, \dots, \lambda_m) = \left(\prod_{k=i+1}^m \lambda_k \right)^{1/(m-i)}$$

$$a(\lambda_{i+1}, \dots, \lambda_m) = \frac{1}{m-i} \sum_{k=i+1}^m \lambda_k$$

The AIC or MDL estimation of the number of sources is taken to be the value of $i \in \{0, \dots, m-1\}$ where $AIC(i)$ or $MDL(i)$ is minimized. Research has shown that MDL is asymptotically consistent, whereas AIC is not consistent, and tends to overestimate the number of sources. However, AIC has been reported to be more robust at a relatively low SNR or with a smaller sample size. Figures 12 and 13 show the AIC (left) and MDL (right) output as a function of i for a simulated object with a single dipole and two dipoles respectively. In our simulations, both of these metrics have achieved near perfect discrimination of single dipole objects versus multiple dipole objects in noise free conditions. Specifically, single dipole objects nearly always had lower minimum AIC and MDL (usually $i = 2$) than two dipole objects (usually $i = 4$ or $i = 5$), so discrimination between one and two dipole objects could be effected solely based on this metric. While performance degrades slightly in noise, performance is still outstanding,

and much better than goodness of fit metrics. Clearly, the remaining task is to assess performance on real data.




Figure 12. AIC output (left) and MDL output (right) for a single object. Minimum shown with red square.

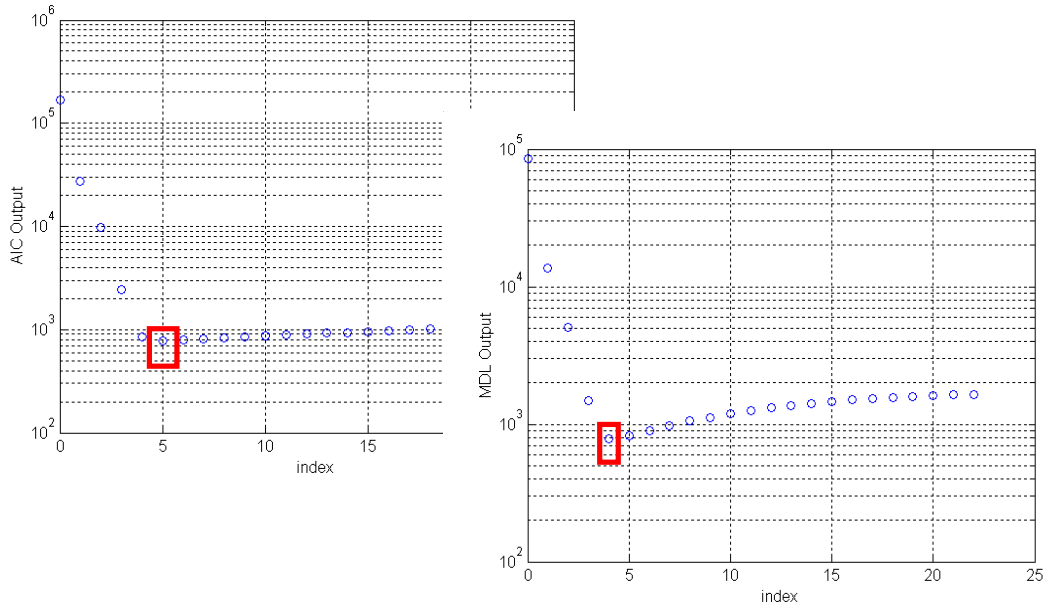


Figure 13. AIC output (left) and MDL output (right) for two objects. Minimum shown with red square.

Signature Separation

As shown in the modeling section, for complex UXO, a multi-dipole model may be necessary to fully replicate the aspect-dependent behavior associated with uncertain target/sensor orientation. A second case in which multiple dipoles may be observed by the sensor occurs when there are two *simple* objects within the view of the sensor. We investigated independent components analysis (ICA) as a mechanism by which to extract the individual targets signatures from sensor data taken at multiple positions over closely spaced objects. ICA assumes that the signatures to be extracted are simple linear combinations of the underlying signals. In the case of multiple closely spaced objects sensed with an EMI system, modeling results suggest that EMI currents present on one object may induce a secondary response in a second closely spaced object, thus invalidating the linearity assumption. As noted above, we utilized our models to assess the relative size of this non-linear effect, which we observed to be small. We also performed several experiments to ascertain whether the effect can be observed in field data.

The first experiment to show proof of concept proceeded as follows. The GEM-3 was used to collect data from two simple (cylindrical and cube) objects. The cylinders were aluminum and steel as were the squares. Cubes were 1" on a side and cylinders were 1" in length and 1" in diameter. Data was taken for the objects in isolation and when they were various distances apart. The sensor was located at several different heights and data was taken spatially. The separation distance was measured from the sides of the objects, not from the center. The distance was equally split from the origin (e.g. the side of each object would be located at +/- 4.5 inches for a 9 inch separation). Data was collected - 16" from the origin to 16" from the origin in 2" increments. Background measurements were made before and after each measurement to correct for background drift. Object separations considered were 3" and 9" and sensor height above the objects was 4" and 7". The data from this collection will also be used to evaluate the signature separation algorithms.

Figure 14 (top) shows the 2790 Hz component measured with the GEM3 for a 9" separation and 4" sensor height. Object 1 is the cylinder and Object 2 is the cube. The blue lines show the quadrature data, the red lines show the in-phase data. Data plotted with the circle and plus symbols show the pattern of spatial response for each of the objects when measured individually. The solid line shows the data measured when both objects are present, and the dashed line shows the data that is predicted by summing the two individually measured signatures. Clearly, linearity does not hold in this situation. Figure 14 (bottom) shows similar data measured with the objects are separated by 9" but the sensor is at a 7" height. In this case, the linearity assumption appears to be valid.

In this experiment, linearity seemed to break down at the 4" height, but was a reasonably good assumption for all other heights. Assuming the target to sensor height distance was such that linearity held, linearity held for all target separations. It is possible that the sensor coil is so close to the targets that the objects distort the transmitted field differently

than one rather than coupling between the sensors, as suggested by Mike Tuley (personal communication). It is also possible that in the near field, linearity does not hold.




Figure 14. GEM response at 2790 Hz (red = in-phase, blue=quadrature) for two objects individually (circle and plus sign), when summed (dashed) and measured with both objects present (solid). Objects are 9 inches apart and 4 inches from the sensor (top) or 9 inches apart and 7 inches from the sensor (bottom).

Next, we began the assessment of algorithms to extract individual signatures from composite measurements. Initially, we considered only simulated data. We considered 5 simulated targets, present in the simulated data either singly, or in pairs. The object characteristics were selected to mimic UXO-like objects, and were estimated using the AETC GEM database, and GEM-3 like object signature were simulated (bandwidth=24 kHz). Specifically, we used data from the database from a 40 mm, 60 mm, 81 mm, M42 submunition, and a 155 mm objects and used the average parameter values estimated from a dipole model fit. Object identifications are: M42 submunition (object 1), 40 mm (object 2), 81 mm (object 3), 155 mm (object 4), and a 60 mm (object 5). These munitions are fairly easy to classify in isolation. The simulations considered a 1 m by 1m grid with samples every 20 cm. For single objects, they were placed at the center of the grid, for paired objects they were placed on the midline of the grid in the y orientation and -25cm and +25 cm in the x orientation. They were assumed to be pointing nose in the y orientation and at a depth from the sensor of 50 cm. Additional simulations were

performed at other orientations and depths and locations, but results were similar so for brevity only these results are included as they are representative.

Classification of the single objects was performed using a traditional Bayesian maximum *a posteriori* classifier, as has been used in our previous SERDP-sponsored research. Performance was studied as a function of SNR and depth. Classification of the pairs of objects was performed by first running a standard ICA algorithm on 9 sets of simulated spatially-collected data, then applying a Bayesian classifier to each extracted signature. Results were only considered correct for the pairs of objects if the classifier identified both objects correctly, i.e. if objects 3 and 5 were present, a classification of 3 and 4 was considered incorrect. Classification in the paired objects simulation was studied as a function of SNR, object depth, and object separation. For a fixed noise variance, confusion matrices were calculated. In this case, SNR is calculated as the ratio between the average object energy and noise variance. At a given spatial location s_i , assuming N in-phase samples in the frequency domain and N quadrature samples in the frequency

domain, energy $E(s_i)$ is defined as $E(s_i) = \sum_{f=1}^N I(f)^2 + \sum_{f=1}^N Q(f)^2$. Average object energy is

defined as the average energy over all spatial locations s_i . This definition of SNR is commonly used in the signal processing literature but has not been shown to be consistent with field values of SNR. Note that since average object energy is utilized over a fairly wide extent, significantly higher SNRs are needed to achieve good discrimination than would be needed if a different energy metric, such as $\max[E(s_i)]$ were utilized instead.

Performance curves for the set of 5 objects presented singly as a function of SNR are shown in Figure 15 (left). Similar curves for a subset of the paired object presentations are shown in Figure 15 (right). Clearly, at a given SNR, the classification of single objects is better than classification of pairs of objects, but it is possible to classify pairs of objects to some degree.




Figure 15. Left panel: Probability of correct classification of 5 objects that appear in isolation as a function of SNR (simulated data). Right panel: Post-ICA classification of 5 objects that appear in pairs as a function of SNR (simulated data). Object

identifications are: M42 submunition (object 1), 40 mm (object 2), 81 mm (object 3), 155 mm (object 4), and a 60 mm (object 5).

Figure 16 shows confusion matrices for the single objects at a fixed SNR of 70 dB (Figure 16 top) and for pairs of objects at the same SNR (Figure 16 bottom). Average percent correct classification is also shown. For this particular combination of object separation, object depth, and SNR, there is an approximately 10% degradation in percent correct classification between the paired object case and the single object case. Also of interest is the performance *enhancement* over Bayesian processing that is performed on the paired object data without performing ICA. The average percent correct performance in this case is 14%, as compared to 85% when ICA is employed. The rationale behind this poor performance without the ICA is associated with the extraction of multiple-dipole information without any prior information. This difficulty has been observed in the multi-dipole fitting described in the modeling section. In a sense, ICA utilized a linear mixing matrix assumption for the spatial data and separates the signatures so that a single dipole feature set can be extracted easily. These results were promising enough that we continued to investigate ICA.

	Object 1	Object 2	Object 3	Object 4	Object 5	Can't be classified
Object 1	94	0	1	2	2	1
Object 2	1	96	0	0	2	1
Object 3	2	0	95	0	1	2
Object 4	1	1	3	94	1	0
Object 5	1	1	1	2	95	0

	o1o2	o1o3	o1o4	o1o5	o2o3	o2o4	o2o5	o3o4	o3o5	o4o5	N/A
o1o2	86	2	0	3	1	1	2	1	1	1	2
o1o3	2	85	3	1	2	2	2	1	1	0	1
o1o4	2	2	85	3	1	3	1	1	2	0	0
o1o5	1	1	1	86	1	3	1	2	2	2	0
o2o3	1	2	3	3	86	1	3	1	2	0	2
o2o4	2	2	3	1	4	83	0	1	1	1	2
o2o5	2	2	3	2	3	1	83	0	0	3	1
o3o4	1	1	3	3	2	1	3	85	0	0	1
o3o5	0	2	2	3	3	0	0	0	88	1	1
o4o5	0	3	1	4	1	1	2	2	1	83	2

Figure 16. Confusion matrices for classifier where objects occur in isolation (top) or in pairs (bottom). Average percent correct for objects in isolation is 95%, and is 85% when occur in pairs when ICA is used, 14% if ICA is not used.

Next, we focused on comparing performance of ICA to other BSS algorithms, since the particular problem we are considering does not meet the independence assumption required by the ICA formulation. We also focused on applying the various approaches to two forms of experimental data: one where UXO objects are measured in isolation and then combined to form mixtures (to test initial proof of concept) and one in which data from closely spaced objects is measured directly and processed. Then, we have considered the situation in which a UXO is present along with *multiple* clutter objects, or in which multiple clutter objects are present without the presence of a UXO. This case was specifically requested following the May 2003 IPR.

In our initial simulations we considered a set of 4 objects which when they occurred in isolation could be discriminated 93% of the time on average. These are the same four objects described above. When these four objects were presented in pairs (with random distances between the objects and placed in the field of view of the sensor at random depths and orientations), the discrimination algorithm could only correctly identify the two objects that were present 22% of the time on average when no BSS algorithm was applied. When ICA was applied prior to discrimination, correct identification occurred 88% of the time on average. Technical details describing the approach and complete results including confusion matrices can be found in the Hu and Collins manuscript published in IEEE TGRS (Hu, W., Tantom, S. L., and Collins, L. M., "Classification of Multiple Closely-Spaced Subsurface Objects: Application of Independent Component Analysis," IEEE Trans. Geosc. Remote Sensing, 42(11), November, 2004, 2544-2554).

In the companion experimental study, four objects were measured in isolation and then combined to create mixtures (objects consisted of a 155 mm Projectile, an M42 Submunition, an Alu Disk (12"), and an Mk118 (Rockeye)). These objects, when presented in isolation, could be discriminated 87% of the time. Failures in discrimination occurred primarily because of poor inversions – essentially local minima in the search space. If known poor inversions were excluded, discrimination performance approached 100%, but this would essentially be an impossible exercise in the real world. ICA-based discrimination of the synthesized two-object complexes was correct 82% of the time. When ICA was not used, both objects were correctly identified 37% of the time. In this analysis we assumed a *single* dipole for each UXO object, which we know to often be untrue, and linearly combined the signatures. Based on these assumptions, performance tracks that obtained for the simulated data. These results are also included in the Hu and Collins manuscript published in IEEE TGRS.

We also considered alternative BSS algorithms. We investigated an approach based on eigenvalue decompositions (EDA), which utilizes second order statistics. The assumptions underlying this approach are more suitable to the UXO discrimination problem using EMI data because independence of the sources are not required. This approach has also been shown in the literature to be more robust when the available data is limited. In the simulation study, average percent correct discrimination performance was 88% using ICA and 91% using EDA. In the experimental data, average percent correct discrimination performance was 82% using ICA and 92.5% using EDA. Based

on these results, it appears that EDA does improve classification performance. Additional performance results using EDA can be found in the Hu and Collins manuscript published in Radio Science (Hu, W. and Collins, L., “Classification of Closely-Spaced Subsurface Objects Using Electromagnetic Induction Data and Blind Source Separation Algorithms”, Radio Science, June, 2004).

Next, we conducted an experiment in which data was measured from two closely-spaced UXO items. For this experiment, a set of 4 UXO targets (M42 bomblet (T1), 40mm projectile (T2), 60mm projectile (T3), and an 81mm projectile (T4)) were measured in various 2 object combinations. The targets were spaced either 1” or 6” apart (measured from side to side), thus each pair of objects occurred twice in the data set. Data was collected with the GEM-3 from -12 to +12 inches in 2-inch increments in a ‘+’ pattern along both the x- and y-axes of the experimental setup, and over 25 frequencies. Data was also collected with each target alone so the extracted independent signatures could be compared to the isolated signatures. Figure 17 illustrates a subset of the individual and extracted signatures, and shows that EDA is doing a fairly good job of extracting the component signatures.

After obtaining the independent signatures from the BSS algorithm, the targets were identified using a Bayesian Classifier. The confusion matrix shown in Figure 18 provides the classification performance. Each pair of two objects was measured in two different spatial configurations, so classification performance can either be 0%, 50% or 100% since both objects must be identified correctly to be deemed a correct classification result. Four out of the six composite mixtures were correctly identified 100% of the time. In one case, the composite was correctly identified once in the two times it was presented, and in the case where it was misidentified, one of the two objects was still identified correctly. In the remaining case, the mixture was never correctly identified, although again, one of the two objects was always correctly identified. These preliminary results from experimental data from overlapping objects are quite encouraging. In the future we would like to apply this approach to data collected in the field.

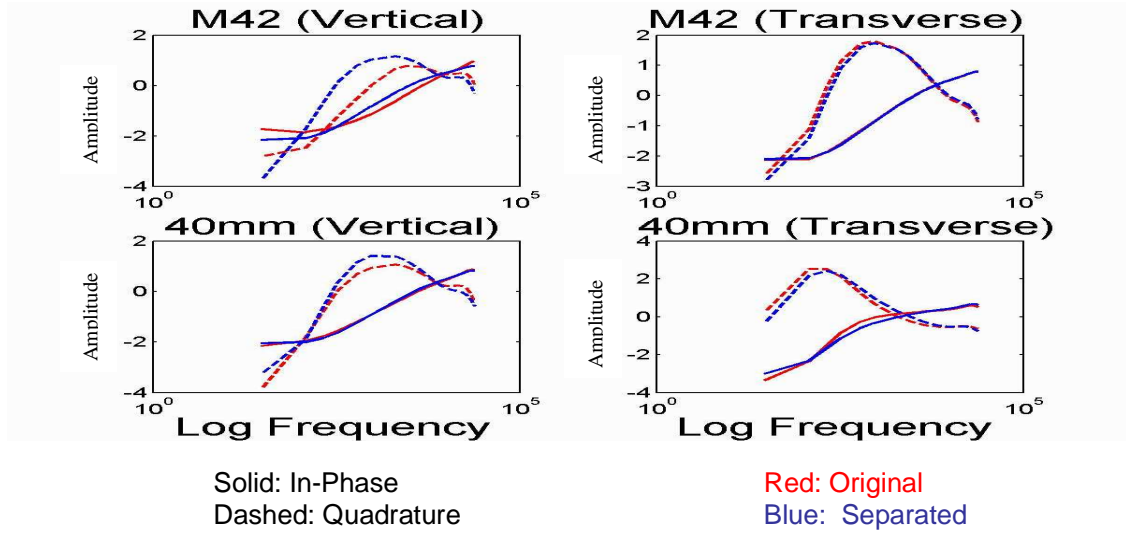


Figure 17. Original and extracted sources obtained for two (M42 and 40mm) of the objects considered in the source separation experiment. Left panels are for the objects oriented vertically, right panel for the objects oriented transverse. The two objects were measured together and separated by 1". Red curves show original sources, blue are the sources extracted using EDA. Solid lines plot in-phase data, dashed lines plot quadrature data.

Target	T12	T13	T14	T23	T24	T34
T12	1.0					
T13		0.5	0	0.5	0	0
T14	0		0		0	1.0
T23		0		1.0		
T24	0		0	0	1.0	0
T34		0				1.0

Figure 18. Confusion matrices listing fractional percent correct identification for experimental data when two targets are present simultaneously in the field of view of the sensor. True targets labeled in rows, classified targets labeled in columns.

In the next experiment, a set of 5 clutter targets (coke can, copper cylinder, lead ring, metallic rock, and a spam can) and 1 60mm UXO target were used to form various target systems with the number of objects present at any given time ranging from 3 to 6. Figure 19 shows an example of one configuration of objects used in the data collection. Data was collected with the GEM-3 from -10 to +10 inches in 2-inch increments along both the x- and y-axes and over 10 frequencies. Data was also collected with each target alone so the extracted independent signatures could be compared to the isolated signatures.

The confusion matrix shown in Figure 20 provides a portion of the classification results using the Bayesian classifier following BSS. These results show that the combined BSS/Bayesian classifier approach was able to identify the UXO correctly (using a library of the four targets from the previous experiment) when supplied a mixed signal consisting of the UXO and several clutter targets. The 60mm was correctly identified in 22 out of 25 trials. In each incorrect case, the candidate item with the second highest value generated by the algorithm was correct. There was no ‘don’t know’ option available to the classifier.



Figure 19. Example of experimental setup for UXO + multiple clutter experiment.

Targets	1 st Choice	2 nd Choice
1234	60mm	M42
1235	60mm	
1236	60mm	M42
1245	60mm	M42
1246	M42	60mm
1256	60mm	81mm
1345	60mm	
1346	M42	60mm
1356	60mm	81mm
1456	60mm	M42

Figure 20. Portion of the experimental results from the UXO + multiple clutter experiment. Targets present are shown in the first column. Classification results for the UXO signature extracted shown in the 2nd and 3rd columns. A 60 mm was always present (target 1).

Area Circumscription

The area circumscription problem was also considered briefly. Figure 21 provides a pictorial representation of the area circumscription problem we are considering in our simulations. In these simulations, we set up contiguous regions in the search space in which anomalies consist of N dipoles, where $N = 1, 2, \dots$. We then apply our prescreeners (goodness of fit, GLRT, MDL, AIC) along transects through the simulation region and plot the output as a function of the distance along the transect. Note that the prescreeners are assuming at least one object – we did not incorporate the ‘no object’ option at this point. We can apply smoothing filters to this output, and then threshold the output to delineate areas of high UXO density. These delineations can then be compared to ground truth. While the results do not match the dipole density exactly since the techniques are estimating signal complexity, not necessarily dipole numbers, our preliminary results are quite promising in that they delineate areas. Figure 22 plots the output of one such simulation along one transect, where the black curve is the output of the AIC prescreener and the red curve shows the smoothed and thresholded plot of UXO density. While these simulated data proved interesting, no field data was available to test the algorithms further. Nonetheless, in a separate effort funded by ESTCP through SIG, we did develop an algorithm that used risk assessment that may also be useful for solving this problem.




Figure 21. Example of the simulation setup for the circumscription problem. Brown area contains no UXO, green area contains widely separated anomalies and dark grey area contains overlapping objects. Cyan lines show simulated transects, data from which are provided to the various algorithms to determine the number of objects present as a function of distance along the transect.




Figure 22. Example results from a single simulation, a single transect. Signal complexity is plotted as a function of distance along the transect. Red curve is a smoothed version of the output.

SERDP/NRL/WES Preliminary Results

Our next focus was on processing the mixed signature data measured for SERDP by NRL, an effort which has been taken over by WES. Initial efforts were primarily limited




Figure 23. A comparison of the performance of classification techniques with and without incorporation of ICA. Classifiers are trained on calibration measurements of isolated UXO. Performance is measured in terms of area under the ROC for which 1 indicates perfect performance and 0.5 indicates random chance.

to processing the GEM-3 data. Results from our initial processing was mixed, and these mixed results prompted several of the simulation studies described below. Essentially, our BSS/ICA techniques worked quite well for all but the largest UXO, particularly when coupled with preprocessing which included the exclusion of ‘background’ data (data with negligible energy from any metal object) from classification. However, as described below, some puzzling results were obtained. In order to investigate whether there were sampling effects, amplitude effects, or signature correlation effects that were causing the degradation in performance, several simulation studies were performed. It should be noted that these simulation studies were based on measured field data as much as possible – simulated data was only used in order to control variables (such as signature correlation) that could not easily be controlled using measured data alone. Our final results on the SERDP/WES/NRL report will be described after the simulation results

Simulations




Figure 24. ROC for each of the four measured objects. Note that the 60 mm object achieves perfect detection at zero probability of false alarm. Only one source is extracted by ICA per object pair.

These results, and additional details regarding the study, are in [Throckmorton, C. S., Tantom, S. L., Tan, Y., and Collins, L. M., “Blind Source Separation for UXO Detection in Highly Cluttered Environments,” *Journal of Applied Geophysics*, 61, 2007, 304-317.]




Figure 25. ROC for each of the four measured objects. Two sources are extracted by ICA per object pair. 60 mm curve follows 81 mm curve where color is ambiguous.

The goal of BSS is to recover independent sources provided only sensor observations that are linear mixtures of independent source signals. ICA is a method for solving the BSS problem. The goal of ICA is to find a linear transformation, \mathbf{h} , of the dependent sensor signals, \mathbf{x} , that makes the transformed sensor measurements as independent as possible.

Therefore, \mathbf{y} is an estimate of the sources. The sources are exactly recovered when \mathbf{h} is the inverse of the mixing matrix, \mathbf{A} , up to a permutation and scale change.

Two key issues in applying ICA are the definition of a measure of independence and the design of algorithms to find the change of basis (or separating matrix), \mathbf{B} , to optimize this




Figure 26. A comparison of the performance of ICA with the original, processed data and the original data with additional pre-processing. Performance is measured in terms of area under the ROC for which 1 indicates perfect performance and 0.5 indicates random chance.

measure. Within a signal processing framework, the ideas of ICA exploit the algebraic structure of higher-order moments of the observed vector, and therefore, the measures of independence are based on fourth-order correlations between the entries of the measured data \mathbf{y} . Several algorithms have been developed, including Jade (Cardoso & Souloumiac, 1993) and FastICA (Hyvarinen & Oja, 1997).

Previous work using independent component analysis (ICA) to classify multiple objects, both in simulation and with measured data, demonstrated the ability to separate and classify each object (Hu et al., 2004). The measured data consisted of four target objects, taken in pairs, with ICA used to extract the signatures of each object. Classification was then accomplished using a library for which sources were generated using a dipole model.

Given the demonstration of the ability of ICA to separate the signatures of two objects, ICA was applied to a field data set (a subset of the SERDP/NRL/WES data set) in which two objects were closely spaced; however, one of the objects was now one of six clutter items. The task was somewhat different from the Hu et al. (2004) study. In that study, all the objects were represented in the library; however, in this task, only calibration data for the UXO were used to determine the presence or absence of UXO in the joint measurements. The performance for a classification algorithm incorporating ICA was first compared to a standard classification algorithm (K nearest neighbors, using a

distance metric). For the ICA-based algorithm, ICA was used to extract sources from the mixtures and these extracted sources were used for classification. For DTKNN, model inversions were performed on the mixtures, and the model parameters providing the best fit to the mixtures were used for feature-based classification. Both algorithms were trained using only the calibration measurements of the UXO in isolation (either as




Figure 27. The four signals used for the toy problem simulation. All correlation coefficients between signals were less than 0.05.

examples of the true sources or as examples of the true model parameters). Utilizing ICA to separate the signatures prior to discrimination provided a clear performance improvement for all but the largest UXO (see Fig. 23). However, the results also had several unexpected elements. First, only extracting a single source for each object pair (either clutter and clutter or UXO and clutter) provided better performance (see Figs. 24 & 25) than extracting two or more per pair. Further, target size was not an indicator of which object would be most easily discriminated. However, the additional pre-processing performed on the data by our lab, which included reducing the size of the data cube considered, did tend to improve performance (see Fig. 26).

Simulations were run in order to gain a greater understanding of the factors that affect ICA, thereby both gaining a greater understanding of the results from analysis of the measured data as well as investigating possible improvements for our analysis methods. Simulations provided an advantage over the measured data since control over the variables for measured data is limited. The following issues were explored through simulation: the importance of the pattern of sampling, how the correlation between signals affects performance, the effect of a mismatch in the number of sources extracted

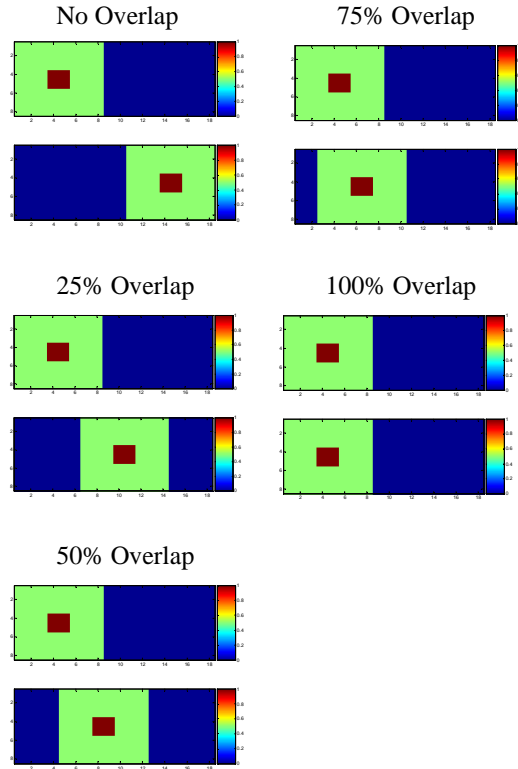


Figure 28. An example of the different levels of overlap for the toy problem simulation. The color bars indicate amplitude, where blue equals zero and red equals 1. The top portion of each subplot represents an object that remains in the same position, while the bottom portion of each subplot represents an object that is brought closer to the first object such that their spatial responses overlap (i.e. their signatures are mixed).

versus those present, and finally the effect of the amplitude pattern (i.e. the weighting pattern that mixes the signals).

Two types of simulations were designed. Although we began with a simulation that incorporated models of UXO signatures and amplitudes patterns, this later proved to offer too little control over variables. So, a second simulation was designed which was more simplistic – a toy problem.

A. UXO Model Simulation

For this simulation, a target was randomly placed within a square meter, and 1-5 clutter items were randomly placed around it at a constant radial distance (separation distance). Thus, the clutter items were all equidistant from the target, but were randomly spaced from one another. The target location was restricted by the separation distance such that clutter could be placed at any radial location and still be located within the square meter.

The target could either be a simulated UXO or clutter; thus, the hypotheses being tested were UXO + clutter versus clutter + clutter.

UXO were modeled as bodies of revolution (BOR). A BOR has two principal coordinates: vertical and transverse. When the sensor is placed along the cylinder axis, only the vertical modes are excited. When the sensor is placed orthogonal to the cylinder axis, only the transverse modes are excited. In reality, the fundamental resonant frequency dominates the sensor response and the corresponding sources can be written as

$$s^v(f) = a^v + \beta^v \frac{f}{f - jf^v}$$

$$s^t(f) = a^t + \beta^t \frac{f}{f - jf^t},$$

where f is the set of measurement frequencies; f^v and f^t are the fundamental resonant frequencies which depend on the material of the object and its geometry (generally, $f^t < f^v$); β^v and β^t modify the strength of the resonances and are determined by




Figure 29. Result of increasing the spacing between samples for the UXO model simulation for 1-5 clutter objects. The performance, measured in area under the ROC, was averaged across separation distance.

target/sensor geometry; and a^v and a^t are constants associated with the EMI response to ferrous objects and are zero for non-ferrous objects. All objects (UXO and clutter) were assumed to be located at a constant depth of 0.5m, however, orientation (φ) and inclination (θ) angles were randomly selected from a range of 0 to π for each object. For the simulation of UXO, a^v and a^t were set to zero, and f^v and f^t were set to 200

Hz and 50 Hz respectively. The resonant frequencies for clutter were randomly selected from a range of 5kHz to 24kHz, with the constraint that $f^i < f^v$.

For these simulations, the area was sampled in an ‘asterisk’ pattern, with the distance between the sample points varying from 5 cm to 25 cm. The sampling patterns were




Figure 30. Results for the case in which clutter no longer has a constant separation distance from the target, i.e. clutter is randomly placed within the square meter without relation to target location. Spacing is increased between samples for the UXO model simulation for 1-5 clutter objects.

centered on the square meter, thus the target might not be centered beneath the measurement pattern. Three separation distances were tested: 5 cm, 10 cm, and 25 cm. The signal to noise ratio (SNR) used in the simulations was 20 dB with respect to the energy of the simulated UXO. A Gaussian detector was designed and tested on the sources extracted by ICA. Performance is measured as area under the ROC where performance of one indicates perfect performance and 0.5 indicates random chance.

B. Toy Problem Simulation

ICA makes the assumption that the signals to be separated are independent (and therefore uncorrelated); however, this is unlikely to be the case for UXO and/or clutter. However, with a toy problem, a set of highly uncorrelated signals can be designed such that issues such as sample spacing and amplitude patterns can be considered separate from the confounding factor of correlation between signals. For this problem, four relatively uncorrelated signals were designed (window pulse train, sawtooth wave, cosine, and square wave) and plotted in Fig. 27. These signals replace the simulated UXO signatures that would normally represent axial and transverse signatures from a dipole model. These signals, being mere constructs to test the function of the ICA algorithm do not have dimensions.




Figure 31. Result of increasing the spacing between samples for the UXO model simulation for 1-5 clutter objects. In this case, an energy detector rather than a generalized Gaussian detector (Fig. 29) was used. The performance, measured as area under the ROC, was averaged across separation distance.

For each signal, an amplitude pattern was designed (varying depending on the problem to be addressed). These amplitude patterns simply represent the weights for each signal that determine the resulting mixed signals. Each pixel can be considered a measurement at a “spatial” location, with the amplitude of the pattern determining the weight of the corresponding signal at that location. The final mixed signal at that location is then the sum of each signal weighted by its amplitude pattern at that location. The amplitude patterns of the first and third signals overlapped completely, as did the second and fourth signal amplitude patterns. The reasoning was to simulate the case where the axial and transverse signals for a single object have amplitude patterns that overlap completely. The overlap between the amplitude patterns of the two signal pairs was then varied to determine the effect of distance between the two ‘objects’ (see Fig. 28). For example, in the top left subplot, at each pixel, the amplitude pattern for one of the objects is zero (blue); thus, the signatures would not be mixed at any location. However, in the other subplots, locations for which the amplitude patterns for both objects are greater than zero would be locations for which the signatures would be mixed based on the corresponding weights for each signal. Performance is measured in terms of correlation between the sources extracted by ICA and the original signals.

Sampling

Two issues were investigated in terms of sampling. First, the effect of sparser sampling on performance was considered. Second, the effect of the pattern of sampling was considered. Both of these issues were investigated in the UXO model simulation.

Fig. 29 shows the results of sparser sampling. The area under the ROC was averaged across target/clutter separation distances. As the number of clutter items increased, performance tended to decrease. Increasing the distance between samples also tended to decrease performance. Theoretically, ICA in a noiseless paradigm with uncorrelated signals requires only a small number of samples of the mixtures (on the order of the number of signatures composing the mixtures) in order to separate the signals perfectly. These simulations, with a high SNR, suggest that for the case of dipole-modeled signatures, the sources extracted by ICA can be used to develop detectors with a high level of performance despite a low level of sampling.

For this simulation, the case for which the clutter location is no longer tied to the target location was also considered. Rather than forcing clutter to be located within a constant radial distance to the target, the clutter was also placed at random within the square meter. The restriction on the target location (that it must be at least the separation




Figure 32. Comparison of performance for the UXO model simulation for three different methods of sampling: asterisk pattern, energy partitioning, and energy threshold. Performance was measured in terms of area under the ROC.

distance from the edge of the square meter) was also removed. As can be seen in Fig. 30, a result similar to Fig. 29 was observed. For high SNR, detectors relying on sources extracted by ICA are fairly robust to decreased sampling rates.

However, this robustness may in part rely on an appropriate choice of detector. For comparison, the effect of sampling rate was considered for an energy detector. In this

case, the decision statistic is based on the energy of the extracted sources. As expected, performance across all cases is much lower with an energy detector (see Fig. 31). In addition to this general performance decrease, however, a lower sampling rate tends to double the drop in performance, e.g. for UXO and 1 clutter, using an energy detector, performance dropped from 0.78 to 0.68 (difference of 0.1), but for the generalized Gaussian detector, performance dropped from 0.99 to 0.96 (difference of 0.03).

For investigating the effect of the pattern of sampling, two energy-based systems were compared to the asterisk measurement pattern. First, a partitioning system was designed in which the energy in the measurement area was partitioned with low energy areas ‘discarded’ (not measured) and high energy areas sampled. Since the partitioning system discarded a large amount of low energy area, it was restricted to 25 samples (rather than the 81 samples used by the asterisk pattern). Another method of sampling areas of high energy set an energy threshold above which all samples were taken, resulting in a variable number of samples. All three measurement patterns are compared in Fig. 32. Little difference in performance was observed for the three methods of sampling. Thus, the high level of performance with detectors utilizing ICA does not appear to be a function of the sampling pattern used.

In the previous simulations, performance with detectors that utilize ICA-extracted sources has been considered with the assumption that high performance is indicative of ICA successfully extracting the original signatures. A simulation to directly assess the success of ICA in extracting sources for different sampling levels was implemented using the toy problem. The four highly uncorrelated signals were used in lieu of modeled axial




Figure 33. The effect of varying the number of samples used for the toy problem simulation. In this case, 0% overlap was used, and the amplitude patterns were modeled from the dipole model, with random orientations. The N highest energy samples were used where N ranges from 800 to 25. Performance was measured in terms of correlation coefficients.

and transverse signatures; however, the amplitude patterns were modeled as the spatial energy patterns generated by the dipole model with orientations randomized across 900 trials. The sensor-to-object separation used to model the amplitude pattern was 0.5 m.




Figure 34. The effect of varying the number of samples used for the toy problem simulation. In this case, 0% overlap was used, and the amplitude patterns were modeled from the dipole model, with random orientations. The N highest energy samples were used where N ranges from 800 to 25. Signal 4 (the square wave) was replaced with another window pulse train that had a correlation coefficient with signal 1 of 0.3. Performance was measured in terms of correlation coefficients.

The signals were mapped to object axial and transverse amplitude patterns as follows: 1st signal – first object axial; 2nd signal – second object axial; 3rd signal – first object transverse; and 4th signal – second object transverse. By using the highly uncorrelated signals, ICA is given a more ideal condition under which to function, and source extraction is not tied to the choice of resonant frequencies for the objects. ICA’s ability to extract the original signals was measured in terms of correlation between the original signals and the extracted sources for different numbers of samples. The samples were chosen as the N highest-energy samples. For the no noise case, ICA was able to extract all four signals near perfectly (correlation coefficients > 0.99). Noise level was increased to 60 dB SNR, and the ability of ICA to extract the signals decreased significantly (see Fig. 33). However, ICA was able to maintain approximately the same level of performance until over 70% of the samples were discarded. These results support the results from the previous simulations that suggest that the performance of ICA-based algorithms is fairly robust to sampling density.

However, these simulations also highlight the decreased ability of ICA to isolate




Figure 35. Example of two amplitude patterns used with the toy problem simulation: simplistic peak with one sample different, and random amplitude patterns.

completely overlapping signals in noise, even under high SNR conditions. Since one signal from each signal pair was more poorly extracted (signal 1 was more poorly extracted than signal 3, and signal 2 was more poorly extracted than signal 4), the




Figure 36. Average correlation coefficients for each signal for the toy problem simulation. The two types of amplitude patterns depicted in Fig. 35 are compared.

question arises as to whether the cause is due to some, albeit small, residual correlation

between the signals is confounding the extraction or whether the relative weights of the signals within a signal pair, in particular the similarity in weights, is causing the detrimental effect. To this end, both the impact of correlation between signals and similarity between amplitude patterns was investigated in the following sections.




Figure 37. Comparison of the performance for different types of amplitude patterns for the toy problem simulation. Dotted lines represent both the random amplitude patterns (filled symbols) and one-sample-different peak amplitude patterns (hollow symbols). The solid lines represent the model-based amplitude patterns.

These simulation results relate to the field data results in several ways. First, these simulations suggest that the performance of an algorithm based on ICA will not necessarily decrease with a decrease in sampling (unless that decrease is substantial). In the field data results, part of the improvement in performance was due to discarding non-informative background samples before performing ICA (approximately 30% of the samples). Since the amount of samples discarded was relatively small, the simulations suggest that performance should not have been affected detrimentally, and by discarding these samples, it is possible that confounding mixtures were excluded from ICA's extraction process.

Also in the field data results, the size of the object did not seem to be correlated with performance (e.g. larger objects were not necessarily easier to detect than smaller objects). Since these simulations suggest that ICA requires only a small subset of samples for high performance, it might be expected that the advantage of large objects (more samples of the UXO signature) does not necessarily occur.

Correlation Between Signals




Figure 38. Comparison of the variability for different types of amplitude patterns for the toy problem simulation. Dotted lines represent both the random amplitude patterns (filled symbols) and one-sample-different peak amplitude patterns (hollow symbols). The solid lines represent the model-based amplitude patterns.

To test the impact of correlation on ICA, the toy problem from the previous section was re-run with signal 4 replaced with another window pulse train (increasing the correlation coefficient between signals 1 and 4 to 0.3). As can be seen in Figure 34, the ability of ICA to extract the other three signals remains the same, but ICA is unable to reliably extract signal 4 for any number of samples. This highlights what is perhaps the greatest limitation of ICA. It is predicated on the assumption that the mixtures comprise independent (uncorrelated) signals. When two signals are correlated, ICA is unable to separate the correlated portions. Thus, one signal will be extracted successfully, containing both the correlated and uncorrelated components; however, because the correlated component cannot be extracted twice, the second extracted signal will only contain its uncorrelated components. Thus, for correlated signals 1 and 4, signal 1 was extracted accurately while only the uncorrelated components of signal 4 were extracted. However, while correlation may affect baseline performance with ICA, the simulations suggest that ICA remains fairly robust to low sampling rates.

These simulations demonstrate that correlation between the individual signatures can have a large impact on the ability of ICA to successfully extract them from the mixtures. Thus, size may be less important in terms of number of samples than it is in terms of the degree of correlation between the UXO signatures and the clutter signatures.




Figure 39. The average of the amplitude patterns producing the 200 highest correlation coefficients for each signal.

Amplitude Patterns

The problem of overlapping UXO differs from a typical ICA application in two regards. First, the signals to be extracted may have more than negligible correlation with each other. The effect of this issue has been considered in the previous section. Second, in typical ICA problems, the signatures that compose the mixtures are assumed to emanate from different locations, resulting in amplitude patterns (weights) that are distinct. However, for each UXO in the problem considered here, the axial and transverse signatures emanate from a single location, thus resulting in amplitude patterns that share similarities that might be confounding to ICA. As an extreme example, consider multiple uncorrelated signals with identical amplitude patterns (equal weighting for all the signals in each sampled mixture). There is no information to indicate that there are multiple signals present rather than a single signal that is the sum of all of the signals, and ICA will not be able to separate the original signals. Similarly, it is possible that having axial and transverse signatures that emanate from the same location (thereby having similar amplitude patterns) could cause a decrease in the ability to separate those signatures.

The toy problem was used to assess the limits of ICA for separating overlapping signatures based on the similarity of the amplitude patterns. Five overlap conditions were investigated with two types of amplitude patterns. The overlap conditions were 0%, 25%, 50%, 75%, and 100%. The amplitude patterns were either random or a simplistic

peak with one sample different for each signal (see Fig. 35), with the random amplitudes changing for each trial but remaining constant across overlap conditions. As in previous toy problems, the amplitude patterns of signals 1 and 3 and signals 2 and 4 overlap completely, representing axial and transverse signatures. As expected, for amplitude




Figure 40. The average of the amplitude patterns producing the 200 lowest correlation coefficients for each signal.

patterns that are almost identical (dashed lines in Fig. 36), ICA fails to extract both of the signatures from each set. Further, the signals that are extracted (3 and 4) are not as correlated with the originals as those extracted from mixtures derived from randomized amplitude patterns. This suggests that similarity in amplitude patterns can have a strong effect on ICA performance, even for highly uncorrelated signals. Proportion of overlap also affects the ability of ICA to extract the four signals; however, performance appears to be fairly robust to increased mixing as long as the signals do not all emanate from the same location (100% overlap).

The results above verify the hypothesis that the amplitude patterns can affect performance, but are not based on physical phenomenology. The performance with dipole-modeled amplitude patterns was compared to the previous simulations in order to assess the degree to which model-based amplitude patterns inhibit successful source extraction. Model amplitudes were substituted for the other amplitude patterns with the orientation of the two ‘objects’ randomized across trials. The results are shown in Fig. 37. In general, performance is as expected. Performance was lower than for the random amplitudes, since the amplitude patterns share some similarities, but better than for the one-sample-different peak patterns. Also, the sensitivity to the proportion of overlap

increased from that observed for the random amplitude patterns. The reason the proportion of overlap affected the three types of amplitude patterns to different degrees was not immediately obvious. However, it was hypothesized that it might be related to the design of each problem. For the one-sample-different problem, the amplitude patterns remained constant from trial to trial. For both the random amplitude patterns and the dipole-modeled patterns, the amplitude patterns changed across trials. While the random patterns were completely randomized (i.e. unlikely to have any one pattern occur more than once), the dipole-modeled amplitudes might share many similarities across trials, dependent on orientation, and these similarities might confound ICA.




Figure 41. Simplistic amplitude patterns generated for the toy problem simulation: peak, ring, 2 peaks, 2 peaks flipped.

Fig. 38 shows the standard deviation of the correlation coefficients for the three types of amplitude patterns. The random and one-sample-different peak patterns have standard deviations that remain below 0.1 while the model-based patterns have higher standard deviations, ranging from 0.15-0.35. The higher variance for the dipole-modeled amplitude patterns is unlikely to be due to orientations being randomized across trials since the random amplitude patterns are also randomized across trial without detrimental effect. Thus, the performance variance for the dipole-modeled patterns was hypothesized to be a function of the amplitude patterns themselves.

To investigate the cause for the high variance with model amplitudes, the amplitude patterns producing the 200 highest and 200 lowest correlation coefficients were averaged, and are plotted for each signal in Figs. 39 and 40. Peaks tended to produce the highest performance while ring-shaped patterns produced the lowest performance for signals 1 and 2. Although the average pattern appears ring-shaped for signals 3 and 4, it is actually the average of patterns that have two peaks rotated about a center point. Thus, the implication appears to be that some amplitude patterns can cause worse performance than

others, and by randomly choosing between amplitude patterns with good performance and those with poor performance, the variance is increased. This was further verified by

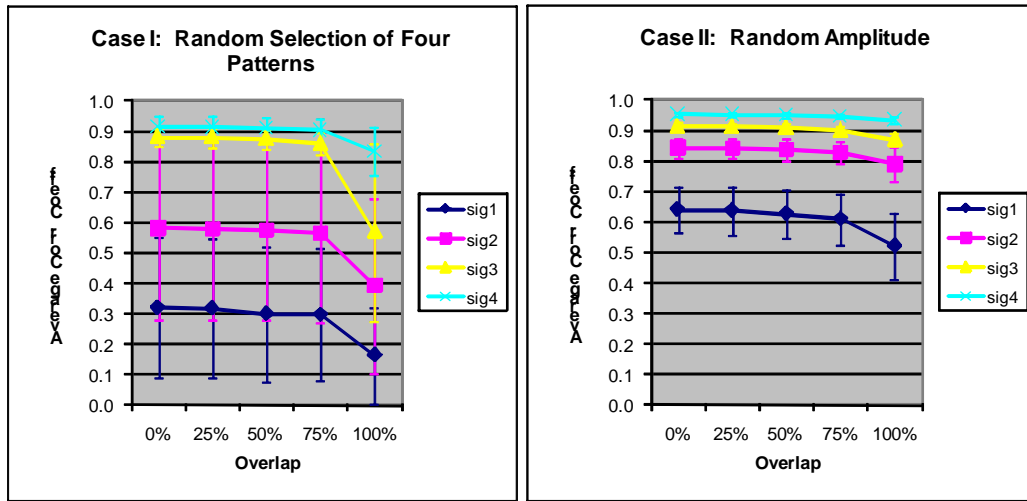


Figure 42. Comparison of performance and variance for two cases of the toy problem simulation: randomly choosing amplitude patterns from a set of four and using random amplitude patterns. Error bars represent one standard deviation from the mean.

simulating several different simplistic amplitude patterns (see Fig. 41) and for each trial, randomly choosing from the set for each signal.

The performance from using this method to select amplitude patterns (case I) is compared to the performance using the random amplitude pattern method (case II) in Figure 42. The mean performance is plotted for each method with the error bars indicating the standard deviation. Two interesting results should be noted. First, the variance for case I is much higher than for case II. Second, proportion of overlap affects case I to a greater degree than for case II. This agrees with the results that were previously noted in Fig. 38. These results may indicate several issues to be aware of with implementation with field data. First, orientation may add variability to results such that performance on field data with one orientation may not be an indicator of performance with other orientations. For an accurate measure of performance, testing on many examples of objects at different orientations may be necessary. In the field data, approximately 75% of the UXO were positioned horizontally. The effect of orientation may be further compounded by the effect of proximity between objects. To fully assess ICA-based algorithms, the examples of objects and orientations will need to be tested at multiple distances.

Number of Sources

The previous simulations suggested some possible reasons for size not being an indicator of performance; however, another unexpected result with the field data was that extracting only one source per object pair produced better performance than extracting

two or more. To investigate this issue, the toy problem simulation was used to investigate the effect of mismatch between the number of sources extracted by ICA and the number of actual signals. The dipole model was again used to provide amplitude patterns which were generated for orientations that were randomized across trials. ICA was used to extract 1-7 sources, with the correct number being 4. Performance was considered for 0% overlap.

The functional form of ICA extracts a number of sources equal to the number of samples




Figure 43. The effect of the number of sources extracted by ICA for the toy problem simulation. In this case, 0% overlap was used, and the amplitude patterns were modeled from the dipole model, with random orientations. The number of sources varied from 1-7, with 4 (the number of signals) being the correct value. Performance was measured in terms of correlation coefficients.

of mixtures; however, a dimension reduction step is typically performed under the assumption that more samples will be measured than there are actual signatures present. Principle Components Analysis (PCA) reduces the data into its basis functions in eigenspace, and typically only those basis functions that represent the majority of the data are retained. This then controls the number of sources extracted by ICA, e.g. if only the one component that most represents the data is retained then ICA will extract only one source. If the number of extracted sources is less than the number of signatures in the mixture, the extracted sources may be successful extractions of some of the original signals, or they may be a mixture of two or more of the original signals. The results in Figure 43 suggest that the latter may be the case when a single source is extracted from the mixture of the four signals. This single source was correlated with all four signals and matched to some extent two of the four (signals 2 and 4). These two were those assigned to a single “object”. When a second source was extracted, the signals from the

second “object” were matched (signals 1 and 3). Performance across all four signals continued to improve (esp. for signals 2 and 4) as the number of extracted sources increased to four, suggesting that some of the mixed sources may have been separated once the number of extracted sources increased; however, the performance increase from separating these mixed sources is not substantial. Other proportions of overlap demonstrated similar results.

These results do suggest that even in the extreme case of extracting one source in the presence of two objects (four signatures), useful information rather than mere noise is contained in that extracted source. Further, these simulations suggest that ICA extracts a source that represents a single object rather than a mixture of all the signatures for both objects. This result may relate to the field data for which performance increased with the extraction of one source rather than two or more. It is possible that with one source, ICA tends to represent the UXO and not the clutter. In the dimension reduction process, the UXO signatures may be represented by a single basis function while the clutter may require more complex representation. This would in turn tend to lower the amount of data represented by each clutter basis function, thus lowering the likelihood that a clutter basis function would be selected during dimension reduction. In this way, ICA may be successfully selecting a UXO source that can then be compared to calibration sources. However, further investigation in to ICA processing will be necessary to fully explain the field data results.

Simulation Summary

Two questions arose from the results with the measured data – why performance might be better with one source extracted rather than more, and why performance did not appear to be dependent on object size. Simulations that varied the number of sources extracted while keeping constant the number of sources present suggested that a single source was capable of representing one of a pair of objects (Fig. 43). This coupled with the results from the field data suggest that the one object represented by the single source is the UXO rather than the clutter. However, further investigation will be necessary to validate this hypothesis.

The simulations of sampling patterns and sampling rates also suggest reasons explaining the lack of dependence of performance on object size. With high SNR, ICA is fairly robust to low sampling rates, hence the advantage of large objects (more samples) may not be applicable. On the other hand, correlation has a large impact on the sources extracted by ICA (see Fig. 34). Thus, the ability to successfully discriminate objects of any size may depend more on the correlation between their signatures and the signatures of clutter than on their size (and the number of samples of their signatures). Furthermore, the amplitude pattern simulations indicate the possibility of a large variance in performance depending on orientation and separation distance. Thus, the performance trends of individual UXO types (e.g. 40 mm) noted in one small data set may not be indicative of the performance results in another data set. The issues driving the high variance observed in the simulation performance should be investigated in more depth through simulations and measured data.

Simulated Time domain EMI signals (EM63)

We consider the problem that there are a total of four objects, two of which are present at a time. In the time-domain, the eddy current time-decay response of a metal target can be expressed as

$$S(t) = \sum_i A_i \exp(-\alpha_i t)$$

where t represents time, A_i is the target amplitude response coefficient and α_i is the i th decay rate. Each of the decay rates corresponds to a natural resonant frequency of the interrogated target and is a function of the target's physical attributes. The time domain EMI sensor response is dominated by the slowest decay rate. The targets of interest in this application (UXO) may be modeled as bodies of revolution (BOR). A BOR has two principal coordinates, axial and transverse, and each coordinate is associated with one dominated decay rate. Thus EMI response of a UXO target can be approximated as a sum of two decaying exponential signals.

$$s(t) = A_v e^{-\alpha_v t} + A_t e^{-\alpha_t t}$$

In the simulation, we assume four objects with decay rates as shown in Table 1. These parameters were chosen to match those of the GEM-3 study so that performance comparisons could be considered. These are based on decay rates extracted from four of the UXO present in the SERDP/NRL/WES overlapping object database as described earlier.

Table 1. Object parameters

Object	α_v	α_t
1	$2\pi \cdot 1000$	$2\pi \cdot 900$
2	$2\pi \cdot 1300$	$2\pi \cdot 1200$
3	$2\pi \cdot 1600$	$2\pi \cdot 1500$
4	$2\pi \cdot 1900$	$2\pi \cdot 1800$

The time signals are linear sampled from 0.04~26 ms in 0.04 ms increments for a total of 650 time samples. A_i values will result from the mixing matrix that is applied.

Comparing classification results from using either a correction matrix method (Hu, Tantum and Collins, 2004) or a template method (Hu and Collins, 2003), we found that under noise free conditions, both methods achieve an 100% correct classification rate, however, in the noisy environment, the template method outperformed the correction matrix method. For example, assuming mixture signals subject to white Gaussian noise with $SNR = 10^{10}$ (estimated from the data set using the same average SNR calculations described previously, e.g. page 15 of this report), Table 2 and Table 3 contain classification results from the two methods. As we can see, the template method achieves an average 63.5% correct classification rate as compared to an average 31.8% correct.

Note that these results are poorer than those obtained with the GEM-3 simulations at the same SNR. This may simply suggest that frequency-domain data is easier to invert than time-domain data.

Table 2: Confusion Matrix Using Correction Matrix Method

		Estimate					
		1&2	1&3	1&4	2&3	2&4	3&4
Truth	1&2	0.86	0	0	0.14	0	0
	1&3	0.92	0.05	0	0.02	0.01	0
	1&4	0.36	0.06	0	0	0.06	0.52
	2&3	0	0	0	0	0	1
	2&4	0.02	0	0	0	0	0.98
	3&4	0	0	0	0	0	1

Table 3: Confusion Matrix Using Template Method

		Estimate					
		1&2	1&3	1&4	2&3	2&4	3&4
Truth	1&2	0.98	0	0	0	0	0.02
	1&3	0.03	0.97	0	0	0	0
	1&4	0	0.01	0.99	0	0	0
	2&3	0.35	0.35	0.19	0.11	0	0
	2&4	0.01	0.01	0.02	0.2	0.76	0
	3&4	0.46	0.32	0.08	0.11	0.03	0

Apply ICA template algorithm to EM63 data

To more accurately assess performance for an EM63 sensor, we model 26 time gates and we only use the first 17 (0.177, 0.191, 0.216, 0.246, 0.286, 0.336, 0.4, 0.48, 0.584, 0.714, 0.883, 1.097, 1.366, 1.714, 2.157, 2.724, 3.445 ms). We simulated the measurements in a ‘+’ pattern (from -10 inches to 10 inches in 5 inch increments along both x and y axes). Thus, there are total 10 spatial measurements.

In this study, we assumed that the number of the objects is two. The simulated overlapping signatures have white Gaussian noise added, and the SNR ranges from 0 dB to 100 dB. The ICA template classification method works perfectly under noise free condition (SNR>100 dB), as observed in the previous simulations. Under the scenario that two UXO targets present, there are a total of six possible UXO pairs. There are also six possible clutter items that can be placed with 4 UXO objects (again, simulating the SERDP/NRL/WES database) and six clutter objects that can be placed in pairs in the context of the simulation. Fig. 44 illustrates that the classification accuracy as a function of SNR, when the ICA template classification algorithm is applied to distinguish the six pairs of UXO. The classification accuracy of 0.167 (around 1/6) implies that ICA template classification algorithm behaves randomly under very noisy conditions (SNR<40 in Fig 44).

Similar classification results can be seen in the other two scenarios: Assuming a UXO-Clutter pair is present, Fig 45 illustrates the classification accuracy of distinguishing 24 pairs of UXO-Clutter under different noise levels; and Fig 46 illustrates the classification accuracy of distinguishing four types of UXO under different noise levels.




Figure 44. UXO-UXO pair classification accuracy using ICA template method vs. different noise levels under the condition that two UXO present.




Figure 45. UXO-Clutter pair classification accuracy using ICA template method vs. different noise levels under the condition that UXO-Clutter pairs present – 2 UXO in library and in data.




Figure 46. UXO Classification Accuracy using ICA template method vs. different noise levels under the condition that UXO-Clutter pairs present – 4 UXO present in library and in data.

As expected, the more UXO in the library, the poorer the performance as the problem is more complicated and uncertainty is higher. While the range of SNR (and the definition of SNR) is not reflective of field standards, the performance trends are consistent with what we would expect with increasing complexity of the problem at hand and of changes in SNR.

EM61 processing

Preliminary work was performed to assess performance of several separation techniques for EM61 processing. One technique is an iterative dipole model fit of the sort proposed by Keiswetter in his 2005 IPR. Another technique is an iterative fit of two dipole models simultaneously. Another fit approach is to use a two dimensional Gaussian shape to model each dipole – this has the advantage of using less parameters to fit the data, thus being less problematic during the inversion process, but has the disadvantage of being non-physics based. The final approach that is being considered is a maximum likelihood estimation technique. An example of a comparison between the iterative dipole model fit and the maximum likelihood fit is shown in Figure 47. The measured data consisted of a 60 mm in the center of the grid with an adjacent 81 mm. A 1 m by 1m square of data with 10 cm sampling was used. It does not appear that the iterative fit (left) is finding a fit for target 2 while the maximum likelihood fit (right) does. While the results are encouraging for the maximum likelihood approach, only limited data has been considered so far.

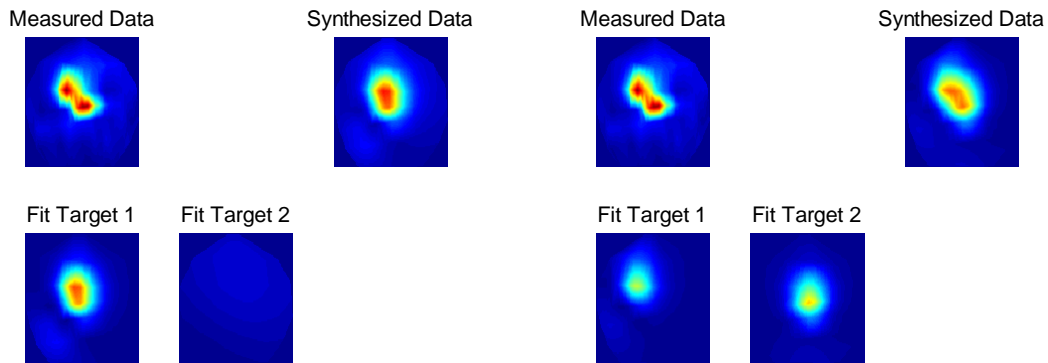


Figure 47. Comparison between iterative fit (left) and maximum likelihood fit (right). In each set of panels, the measured (field) data is shown on the top left, the two individual fits to the data are shown on the bottom left, and the synthesized data from the two individual fits is shown on the top right.

From the simulations we learned that ICA is a promising technique for source separation, but for the UXO it can be impacted by (1) correlation between either the sources or their amplitude patterns; (2) the number of sources extracted versus the number of sources present; (3) mismatched energy between the sources; (4) overlap between the modes of the UXO and the clutter. It was also noted that complicated sampling strategies are not necessary – simply utilizing the highest energy samples in a spatial area is sufficient.

SERDP/WES/NRL data redux

The data set consisted of 18 single object, 67 dual object, and 1 no-object (hole) measurements. There were four UXO target types (105 mm, 81 mm, 60 mm, and 40 mm) and six clutter items. The single object measurements were provided for calibration and consisted of one measurement of each UXO target type in three different positions: horizontal, vertical nose up, and vertical nose down. In addition to those 12 measurements, each clutter item was also measured once for calibration purposes.

The dual object measurements consisted of 14 measurements of the 105 mm, 60 mm, and 40 mm targets; 19 measurements of the 81 mm target; and 6 measurements of combinations of clutter items. The majority of the target measurements (46) were of horizontal targets, with the 81 mm and 60 mm also measured once in a vertical nose up and once in a vertical nose down position. One of the positions for an 81 mm target measurement was not listed, and ten of the 40 mm target measurements were in the vertical nose up position. The first four clutter items (or fragments thereof) were paired with UXO targets 11-15 times, and the fifth and sixth were paired with UXO targets 4 and 3 times respectively.

As mentioned previously, we generally expected that a bigger object would be easier to classify than a small object since there would presumably be more samples with the response to a big object, but the 105 was the worst performer *when analyzed in the same manner as the rest of the objects*. Classification was better for the 105 mm when ICA




Figure 48. Area under the curve for classification of each ordnance type from the SERDP/NRL/WES database. AUC is plotted as a function of the number of grid points used to perform ICA. Training data for the classifier was obtained from the isolated UXO items.

was performed on the mixed signals when the data was arranged as spatial data x frequency; whereas for the other objects the data was generally arranged freq x spatial data. For example, see the results provided in Figure 48. Performance is plotted versus the number of sample points using a radial zoom pattern. The radial zoom patterns is a method for selecting data points for processing. A circle is centered on the target and all sample points within that circle are used for processing. By varying the radius of the circle, the amount of data used for processing can be varied and its impact assessed.

Performance for the 40 mm, 60 mm, and 81 mm UXO is similar; however, performance for the 105 mm UXO, when analyzed in the same way, is quite poor. However, performance can be significantly improved by transposing the mixed signals, increasing performance to the same level as the other objects. It is unclear at this point the basis for this result.

Generally, if ICA is not performed on the data, classification performance is chance, i.e. 25%. This suggests that given single object classification performance, ICA followed by a Bayesian classifier improves performance to around 75% correct.

V. Conclusions

Simulations, test stand data, and field data indicate that ICA/BSS techniques can be used to extract individual ordnance signals from mixtures measured by EMI systems when UXO and clutter are in close proximity. The EDA approach, which is less sensitive to correlated objects, provided the best performance of the approaches considered. While

there are several caveats, and research remains to be done, this study has demonstrated that ICA can restore classification performance from essentially chance levels when objects are closely spaced. Issues related to determining when there are multiple objects present appear most pressing, and may require more densely sampled data. Test stand results were promising, but a conclusive test on a realistic data set is necessary prior to any firm conclusions.

VI. Technology Transfer

Source code for multiple dipole model provided to ERDC and AETC.

VII. Manuscripts Published and Presentations Given

Hu, W., Tantom, S. L., and Collins, L. M., "Classification of Multiple Closely-Spaced Subsurface Objects: Application of Independent Component Analysis," IEEE Trans. Geosc. Remote Sensing, 42(11), November, 2004, 2544-2554.

Hu, W. and Collins, L., "Classification of Closely-Spaced Subsurface Objects Using Electromagnetic Induction Data and Blind Source Separation Algorithms", Radio Science, June, 2004.

Zhang, Y., Collins, L. M., Yu, H., Baum, C., and Carin, L., "Sensing of Unexploded Ordnance with Magnetometer and Induction Data: Theory and Signal Processing", IEEE Trans. Geosc. Remote Sensing, Vol. 41 (5), pp. 1005-1015, May, 2003.

Zhang, Y., Collins, L., and Carin, L., "Unexploded Ordnance Detection Using Bayesian Physics-Based Data Fusion", Integrated Computer-Aided Engineering, Vol 10, pp. 231-237, July, 2003.

Throckmorton, C. S., Tantom, S. L., Tan, Y., and Collins, L. M., "Blind Source Separation for UXO Detection in Highly Cluttered Environments," Journal of Applied Geophysics.

Zhang, Y., Collins, L. M., and Carin, L., "Unexploded Ordnance Detection and Discrimination via Model-Based Statistical Sensor Fusion," UXO/Countermining Forum, Orlando, Florida, September, 2002.

Miller, J., Keiswetter, D., Collins, L., and Torrone, P., "UXO Variability: A Fundamental Limit On Discrimination Performance", UXO/Countermining Forum, Orlando, Florida, September, 2002.

Collins, L., Zhang, Y., Hu, W., and Carin, L., "Statistical Signal Processing For Uxo Discrimination: Performance Results For Isolated And Overlapping Objects", SERDP/ESTCP Symposium, Washington, DC, December, 2002.

Zhang, Collins, and Carin, "Model-based Statistical Signal Processing for UXO Discrimination: Performance Results for the JPG-V Demonstration", Detection and Remediation Technologies for Mines and Minelike Targets VIII Conference, 2000 International Symposium on Aerospace/Defense Sensing and Controls, Orlando, Florida, April, 2003.

Zhang, Collins, and Carin, "Physics Model Based Unexploded Ordnance Discrimination Using Wideband EMI Data: Can More Complex Models Help Discrimination?", Detection and Remediation Technologies for Mines and Minelike Targets VIII Conference, 2000 International Symposium on Aerospace/Defense Sensing and Controls, Orlando, Florida, April, 2003.

Hu, Wang, Schofield, and Collins, "Separation of Overlapping Signatures in EMI Data", Detection and Remediation Technologies for Mines and Minelike Targets VIII Conference, 2000 International Symposium on Aerospace/Defense Sensing and Controls, Orlando, Florida, April, 2003.

Collins, L., Hu, W., and Schofield, D., "Statistical Signal Processing for UXO Discrimination: Performance Results for Isolated and Overlapping Objects," EUDEM-2, Brussels, Belgium, September, 2003.

Collins, L., Hu, W., Schofield, D., and Carin, L., "Statistical Signal Processing For UXO Discrimination: Performance Results For Isolated And Overlapping Objects", SERDP/ESTCP Symposium, Washington, DC, December, 2003.

Collins, L. M., Throckmorton, C. S., Tantum, S. L., and Carin, L. "Statistical Signal Processing For Uxo Discrimination: Experimental Performance Results For Overlapping Objects, 2005 SERDP/ESTCP Symposium.

Tan, Y., Tantum, S. L., and Collins, L. M., "UXO Discrimination Using Blind Source Separation," Proc. SAGEEP, Atlanta, GA, April, 2005.

Wang, C. and Collins, L. M., "Feature selection for physics model based object discrimination," Proc. Detection and Remediation Technologies for Mines and Minelike Targets X Conference, 2005 International Symposium on Aerospace/Defense Sensing and Controls, Orlando, Florida, March, 2005.